Understanding the U-Shaped Curve: Central Claims and Applications for AI*

B. Briskey
Indiana University-Purdue University
Fort Wayne
brisbm02@students.ipfw.edu

M. Greidanus Romaneli
Indiana University-Purdue University
Fort Wayne
mgreidan@purdue.edu

D. Hale
Indiana University-Purdue University
Fort Wayne
haledt01@students.ipfw.edu

J. Licato
Indiana University-Purdue University
Fort Wayne
licatoj@ipfw.edu

ABSTRACT

Representational Redescription theory as proposed by Karmiloff-Smith [1992] investigates the changes in behavioral performance and "level of representation" as children experience new domains of knowledge. In order to introduce the applicability of this theory's propositions to knowledge acquisition in both children and Artificial Intelligence systems, we analyzed the experimental literature in Representational Redescription, as well as in the closely related theory of Neuroconstructivism.

1 INTRODUCTION

Investigations into how children process information have typically focused on specific factors, including biological and socio-cultural constraints, environmental cues, and innate predispositions to attend to certain stimuli. All of these factors have demonstrated important influences on how children learn and communicate new ideas and abilities. However, much of the evidence for these effects has come from observations and interpretations of behaviors. While behaviors may be readily observed and interpreted, processes taking place within the child's mind that may be influencing behavior may be quite difficult to identify. One theoretical perspective that attempts to explain these internal processes, named "representational redescription theory," seeks to explain how children can acquire representations of their external and internal environments that attain increasing levels of complexity and flexibility [Karmiloff-Smith, 1992]. Insight into this process may also have important applications for machine learning and AI, by facilitating the development of progressively more complex skills and capacities in intelligent agents. Yet, such applications can only exist if the concept of 'representation' and its change throughout development is coherent. Unfortunately, work that refers to representational redescription theory has not been consistent on the point of what it means for a representation to change.

ACM Reference format:
B. Briskey, M. Greidanus Romaneli, D. Hale, and J. Licato. 2017. Understanding the U-Shaped Curve: Central Claims and applications for AI. In *Proceedings of Modern Artificial Intelligence and Cognitive Science Conference, Fort Wayne, Indiana USA, March 2017 (MAICS'17)*, 8 pages.

We will address the confusion over what a representation is and how it changes, as predicted by representational redescription theory. This will be followed by an account of Karmiloff-Smith's original operationalizations of key concepts of the developmental curve, based on her landmark 1992 publication. We will then discuss subsequent literature on these predictions, while also clarifying the lack of consensus on how to test representational change independently from behavioral mastery. Finally, we will present alternative methods for measuring this change, in specific neuroconstructivist computational modeling, as well as possible applications of representational redescription to artificial intelligence.

2 OVERVIEW OF REPRESENTATIONAL REDESCRIPTION

Representational redescription theory, originally developed by Annette Karmiloff-Smith, attempts to explain how a child represents the external environment within their mind, changes these representations through continued interaction with the environment, and eventually reaches a higher degree of both behavioral mastery and metacognition. Karmiloff-Smith [1992] divides this process of change into three phases, characterized by two factors: the level of performance on a task, and the 'level of representation[al development],' an abstract measure whose lack of a definition is the cause of much confusion. The relation between levels of performance and representational development form what Karmiloff-Smith calls a "U-shaped curve" (Figure 1). In phase 1, performance is high but level of representation is low; the child does not possess any metacognition about the performance of that specific behavior, but he or she has learned to associate the behavior with its context through observation and feedback from the environment. As the behavior is repeated, they begin to develop a higher level of representation marked by, for instance, theory-building and generalizations. In phase 2, the level of representation within a given micro-domain goes up and is associated with lower

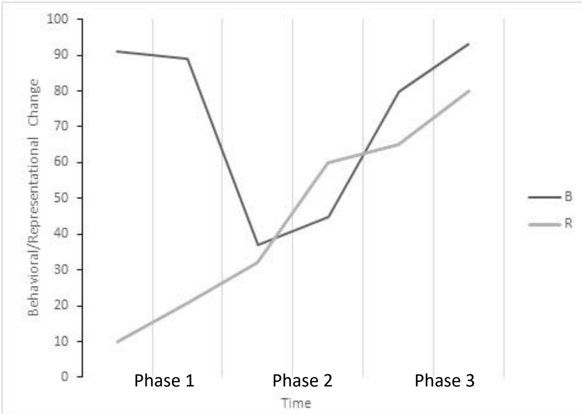

Figure 1: The theorized U-shaped curve of behavioral performance and level of redescription as a function of time, as originally presented in Karmiloff-Smith [1992]. Note the somewhat ambiguously labeled axes.

performance as a result of a higher disregard for external feedback and an increased focus on internal processes.

Finally, as the process of representational redescription continues to phase 3, the child demonstrates more representational flexibility by not only being able to reflect on and describe the rationales for their behavior, but also through accounting for exceptions to the pre-established representations. Thus, the developmental graph of a specific microdomain shows a linear development of level of representation, but at the same time a non-linear development of performance, reaching its lowest level at phase 2. A graph of the U-Shaped curve, equivalent to the one initially proposed by Karmiloff-Smith [1992], can be seen in Figure 1. This U-Shaped curve, originally presented in Karmiloff-Smith [1992] as a hypothetical model, has room for interpretation. Other authors, such as Pinker [1995], have referenced this decrease in behavior performance in other terminology.

Initial support for the change in the level of representation relied specifically on grammatical development, with changes in representation marked by the ability to analyze grammatical elements on a pronoun-noun phrase in French [Karmiloff-Smith, 1979, 1986] and the ability of separating the grammatical elements of American Sign Language (ASL) [Newport 1981]. The inference that representational change occurred in these cases stemmed primarily from the same evidence that suggests a decrease in behavioral performance.[1] In those studies, markers of less accurate behavior (e.g., breaking down a pronoun into a more complex and incorrect phase, breaking down a sign in ASL instead of delivering it fluidly) were indications that children were hyper-aware of the grammatical elements in their utterances/gestures, to the extent of prioritizing the acknowledgment of these internal processes over what was initially correct performance. Late-occurring behavioral performance errors in both examples constitute a large part of the evidence for representational change, however, Karmiloff-Smith further analyzes specific aspects of those mistakes in order to support the U-shaped curve and representational redescription. These aspects, as presented in Karmiloff-Smith [1992], were originally described as accompanying an increase in the level of complexity of representations, demonstrated by the presentation over time of abilities not previously available. They include the abilities to (1) analyze a procedure/representation into its meaningful parts, (2) identify and understand relationships between elements of a larger whole (e.g., morphemes within words), (3a) form general theories about a specific procedure or micro-domain, as well as to (3b) later correctly identify and address exceptions to such theories, (4) contrast characteristics or functions of similar elements (e.g., definite vs. indefinite articles) and (5) elaborate verbal explanations of rationales for either correct or incorrect behavior.

These specific markers of representational change, however, were not as clearly operationalized as the decrease in behavioral performance itself in supporting the U-shaped curve, and may be to blame for some inconsistencies in the research following the original publication of the theory. The confusion about the meaning of the U-shaped curve and the ambiguity of representational redescription's central claims have led to a variety of incompatible interpretations. We will now summarize these interpretations.

3 WORK ON REPRESENTATIONAL REDESCRIPTION

3.1 Language

One of the domains explored in Karmiloff-Smith [1992] was language, with an emphasis on semantic and morphosyntactic development. Critten, Pine, and Steffler [2007] are early proponents of the application of the theory to other language related microdomains, such as spelling. In their 2007 study, the lowest level of representational redescription in children's spelling development was operationalized as inflexible behavioral mastery without conscious access to the knowledge in the microdomain, while the second and third phases were characterized, respectively, by overgeneralization errors progressing to less performance errors without full rule understanding, and finally proficient task performance with complete comprehension of applicable spelling rules. The researchers were able to demonstrate a behavioral developmental curve, in which the decrease in performance was associated to an over-reliance in specific phonetic or morphological spelling strategies.

Expanding on this first study, Critten, Pine, and Messer [2013] continued to analyze representational and behavioral change within

[1] Respectively, the breakdown of the correct utterance "mes voitures" [1st person + plural possessive "my" + plural "cars"] into incorrect forms in which grammatical person and number were separated, "les miennes des voitures" [Karmiloff-Smith 1979, 1986] and the breakdown of initially "holistic" ASL signs into "staccato-like" movements, in which the parts of the sign were broken down based on its morphological markers.

2

the microdomains of spelling recognition and production. In the spelling recognition task, children were asked to choose between one correct and two incorrect spellings of a list of words, including regular and irregular verbs, as well as non-verbs. Performance and ability to justify their choices, as well as why the other alternatives were incorrect, were taken into account when assigning phases of representational development to children. While children in the second-lowest phase had high accuracy in choosing the correctly spelled option (following a phase of low accuracy), subsequent phases demonstrated lower accuracy due to the emergence of new mistakes (e.g., identifying [-ed] as a marker of past-tense in verbs and rejecting correct spellings of irregular verbs for lacking the suffix). Finally, in the last phase, children not only regained high performance, but were also able to explain verbally the rationale for their choices. This pattern of development, composed of a decrease in behavioral performance (marked by accuracy) due to the over-regularization of newly identified elements of language (e.g., [-ed] suffix marking past tense), is in accordance with the original propositions of the U-Shaped curve and the patterns of behavioral and representational development proposed by Karmiloff-Smith [1992].

In a similar spelling production task, the researchers presented children with a spelled word or pseudoword, and an assortment of letter magnets. Children then had to use the letter magnets in order to transform the previous word into a new word proposed by the researcher. Performance was measured in terms of their ability to accurately create the new word and verbally express their justification for both how they transformed the word and how the words differed. Successful performance of the task involved replacing the first letter of the word or pseudoword (always a consonant) with a different consonant. In this second task, providing an incorrect explanation coupled with lower task accuracy was considered by the researchers as a marker of an intermediate phase, which would correspond to phase 2 of the U-shaped curve. This spelling production task, unlike other literature on Representational Redescription and spelling development, did not directly address morphological development, while the skills tested were based purely on phonetic development and letter recognition.

Critten et al. [2007] and Critten et al. [2013] utilized children's verbalizations in order to establish phase distinctions that are not present in representational redescription theory, creating what would be a new section of phase 2 of the U-shaped curve, separating between phonological and morphological overgeneralizations and errors. This distinction appears to be specific to the spelling microdomain, and may not necessarily extend to other developmental microdomains. In proposing the subdivision, however, the researchers still relied heavily on verbal explanations, even in the phases in which children are predicted to not be able to produce them. Evidence based on children's explanations, although effective in distinguishing between implicit (phase 1 in the U-shaped curve, when children are only procedurally aware of the task) and explicit representations (phase 3 in the U-shaped curve, when children have a higher level of representation and are able to verbalize rationales for their behavior), is not sensitive to the markers of representational change in intermediary phases.

Critten, Sheriston, and Man [2016] further examined the applications of representational redescription theory in the microdomains of spelling recognition and production. Two groups of children from three different UK schools, from year 1 and year 2 respectively, performed two separate tasks. Much like Critten, Pine, and Steffler [2007], children were presented with three words, one spelled correctly and two incorrectly, and were prompted to give an explanation as to why their choice was correct and why the other two choices were not correct. In addition, children were asked to spell words that had been previously presented, although in a different order than the child's previous trial. Critten et al. [2016] operationalized representations as children's explicit explanations about spelling performance using the same evaluation criteria as Critten, Pine, and Steffler [2007]. The participants were then grouped together based on an overall level of representational development, determined by the type of explanation given for their performance on the tasks. Children in the group corresponding to phase 2 often made overgeneralization errors and therefore demonstrated representational inflexibility. The researchers maintained the subdivisions proposed in previous studies [Critten et al. 2007; Critten et al. 2013] and provided additional empirical support to a division between morphological and phonetic errors in phase 2. This categorization of groups supports the notion of the U-shaped developmental curve in representational redescription as children's explanations became more advanced as well as having better performance on spelling recognition tasks once their spelling knowledge became more flexible.

Lorandi and Karmiloff-Smith [2012], on the other hand, analyzed children's morphological knowledge in Brazilian Portuguese. They did so through recording the spontaneous occurrence of variant forms of verbs (e.g., overgeneralizations and neologisms inconsistent with the standard norm of the language), as well as through a morphological test in which children were given a nonsense grammatical base (created in accordance with common word structures in the language) and a semantic context, and were prompted to produce inflected forms of such words. While measuring the incidence of errors in these tasks, the researchers found support for an increase in performance and for children's awareness of morphological markers of location, agency, and tense, among others. Although the results in this study do not make evident a U-shaped curve of behavioral performance, they support an increase in level of representation in accordance with the claims of representational redescription theory (e.g., in children's ability to identify and utilize morphological markers in novel productions).

Evidence from microdomains within language, as a whole, indicate the existence of a U-shaped curve very similar (if not equivalent) to the one proposed in representational redescription theory. Still, the simultaneous development of some of these microdomains (such as phonetics, morphology, and letter recognition in the spelling tasks described) and the reliance on verbal explanations provide confounding factors. Thus, these methodologies, although shedding some light into the relationship between representational redescription and language development, do not offer generalizable, agreed-upon insights into this process.

3.2 Mathematics

While representational redescription is argued to be domain-specific — that is, representations in each micro-domain undergo the developmental process independently — the investigations seeking to support or refute the theory's claims vary greatly in terms of domains of knowledge analyzed. One of the fields of experimental research mentioned by Karmiloff-Smith [1992] was arithmetic. Voutsina [2012] examined the presence of markers of representational redescription through testing children's ability to identify number pairs, i.e. what numbers can be added together to obtain a specified number. A qualitative analysis of ten 5-6 year old children's strategies in finding all pairs of numbers revealed that children's explanations of behavioral strategies changed over the course of the study, based on a set number of phases determined by the researchers. Children in the first phase were able to distinctly view the procedures for solving each step; this first step involved children only trying to solve what number pairs made up the number bond without regard to linking the different number pairs together. Later, in the second and third phases, children were able to manipulate and link different aspects of knowledge. After observing how they solved the problem, children developed a strategy for organizing number pairs for number bonds. For example, for the number bond "9", children would replace [0+9] for [9+0], in a strategy called swapping, rather than randomly organizing number bonds, such as [2+7] and then [4+5].

Thus, Voutsina [2012] provides evidence for a constant increase in level of redescription, marked by the emergence of new abilities and strategies from phase to phase. Furthermore, throughout the study, strategies used by the children (such as the ordering strategy) resulted in constant behavioral performance. Considering also that even in the first level of representation, described by the researchers, there were already signs of a higher level of redescription; Voutsina [2012] provides evidence for what appears to be the upwards slope of both behavioral performance and level of redescription. It does not, however, provide evidence supporting or refuting the decrease in performance characteristic of the U-shaped curve, as there was no operationalization for verifying the process of spontaneous redescription of knowledge.

Likewise, Simpson and Stehlikova [2006] examined representational redescription in the domain of Algebra. Rather than with children, this study observed a college student's development across a span of three-years in a case study. Initially, the participant followed strict procedures in order to solve the problems, such as solely adding up the numbers in the problems. The "second shift of attention," or the first explicit level, was operationalized as using the knowledge gained from the familiar procedure towards similar problems that would also use a similar procedure. During the study, as the participant gained a more in-depth understanding of the representation, she made errors; eventually the participant realized the errors and corrected them. As time went on, she verbalized the relationships between some of the numbers before moving to the last phase of representational redescription where she utilized her knowledge about the problems and procedures the professor gave to similar, extra problems and the operations needed to solve them. This participant's performance during the study supports the existence of the U-shaped developmental curve in representational redescription theory. Despite this support, there was no operationalization for verifying the process of spontaneous redescription of knowledge; Simpson and Stehlikova [2006] acknowledged that there was no spontaneous redescription of knowledge, as the participant did not suddenly gain an insight to utilize a specific abstract algebra strategy (z-subtraction) even though she realized and corrected her errors. Even though there was no spontaneous redescription of knowledge, support for the U-shaped developmental curve in Simpson and Stehlikova [2006] was found as the participant's performance at first decreased before increasing.

3.3 Drawing

The domain of drawing involves representations; these representations can be modified and applied to various drawing microdomains. In Picard and Vinter [2006], it was hypothesized that exposure to a model would increase performance in an associated task. Children were tasked with drawing pictures with components that either had no separation or separation by two parts or multiple separations. It was discovered that breaking down knowledge helped change the representation from implicit to explicit; through this decomposition process redescription could occur. During decomposition, children made new knowledge connections such as being able to identity and draw the different parts of the drawing, indicative of knowledge being redescribed.

Similarly to Picard and Vinter [2006], Hollis and Low [2005] investigated representational redescription in the drawing microdomain. The purpose was to investigate environmental constraints that impacted the flexibility of representational redescription in the drawing domain. The study had 315 children, whose ages ranged from 6 to 9 years old. The participants were tasked with drawing pictures of a typical house and person before being tasked with drawing a picture of a "pretend person."

Afterwards, researchers asked the participants to talk about their "pretend person" and later put the children into either the distraction condition, draw alone condition, or the explanation about why this is a "pretend person" condition.

The participants were tested a total of four times. The results from Hollis and Low [2005] support representational redescription. It was observed that children aged 6-7 were at the beginning phase of redescription; their procedures were inflexible. This changed over time, however, when their knowledge became more flexible after exposure to the explanation condition; this condition was associated with a greater rate drawing modification during the middle of the drawing process. This change in knowledge, from inflexible performance with few mistakes, to flexible performance with many errors, is in support of the U-shaped developmental curve.

3.4 Block Balancing Tasks

Similar to the language domain, representational redescription has been widely studied in the block balancing microdomain, where children are tasked with balancing blocks on a beam. Messer, Pine, and Butler [2008] hypothesized that representations are similar in a domain, even if the tasks vary. Children in an equivalent of phase two of representational development tended to have lower performance than children at later representational levels, even though they used consistent strategies to try to solve the block balancing task. Although the children utilized the same strategies in order to solve the block balancing task, the children who had not yet redescribed their knowledge had more errors while solving the task, in comparison to children who redescribed some of the knowledge. This is in contrast to children at the first phase who could not explain why their behavior on a task but still had high behavioral performance. This change in knowledge, from inflexible performance with few mistakes, to flexible performance with many errors, along with the decrease in performance for children at the abstraction level, supports the U-shaped developmental curve.

Representational change was operationalized as a difference in performance in tasks. Cheung and Wong [2011] hypothesized that participants would be unaware that they transferred the strategy of a geometric-center theory for balancing the blocks to different tasks. Behavioral performance was operationalized by how many times children successfully placed blocks in the middle of the beam, in the middle before moving them, and only in the center area of the beam for each block trial. Eight of the twelve participants utilized geometric-center theory across the task; this theory states that children believe that all objects can be balanced at the center. Despite utilizing this strategy, children were unaware of this knowledge, supporting the hypothesis; they tried to complete the task but were unable to explain what their strategy utilized. Although the U-shaped pattern did not appear in the study, Cheung and Wong [2011] argued that the absence of the U-shaped curve does not discredit the representational redescription model, as it is not the main premise behind representational redescription. Other researchers, however, have utilized this absence of a U-Shaped curve in their data as a basis for refuting Representational Redescription theory [Krist, Horz, and Schonfeld, 2005].

Krist, Horz, and Schonfeld [2005] disputed representational redescription in the block balancing task due to the absence of the U-shaped developmental curve. The U-shaped curve was operationally defined as children utilizing the geometric-centric theory during the block balancing task. The study tasked sixty five children, between 4 and 8 years old, with balancing wooden blocks, evenly and unevenly weighted, on a beam. In order to not influence the theories that the children made while completing the task, researchers did not prompt children with questions such as if some blocks are "unbalanceable." Rather, the researchers asked the children why/how blocks can be balanced. It was predicted that success with the task would follow the U-shaped curve with the unevenly weighted blocks: a decrease in performance subsequently followed by an increase due to the redescription of block knowledge from correct block readjustments. This change in performance would be indicative of the children's representational change from the first explicit level (E1) to the later levels (E2 and E3). Success with evenly weighted blocks, however, was predicted to either stay the same or increase with children's age.

Results in Krist et al. [2005], however, showed no support for the U-shaped curve. While performance success improved with evenly weighted blocks, in unevenly weighted blocks performance success decreased along with children's correct block rearrangements. Kris et al. [2005] explained this lack of evidence due to differences in the methodology utilized in Karmiloff-Smith's *On Modularity*; Krist et al. [2005] had more participants, a standardized system to measure children's performance, and did not prompt the participants to make predictions but rather why/how blocks can be balanced in contrast to the original study. Additionally, these results differed due to the fact that the theory of representational redescription has not made clear an operationalization for verifying the process of spontaneous redescription of knowledge, which resulted in these studies using diverging, and often conflicting, methodologies. Despite criticisms and lacking evidence for the U-shaped curve, Krist et al. [2005] advocated for finding domains representational redescription can occur in, rather than throwing out representational redescription theory. Although skeptical, Krist et al. [2005] acknowledged a possible connection between representational redescription and "modern connectionism with developmental neuropsychology," or neuroconstructivism.

3.5 Neuroconstructivism

Achieving a consensus regarding what constitutes a "representation," along with how such a representation might adapt or change, appears to be a critical problem when testing representational redescription. In the studies previously described, behavior has been linked with encoded information states, i.e. representations, effectively associating changes in behavior with changes of these representations. In addition, verbal reports of one's understanding can lead to alterations of this understanding, and only the representations that can be verbalized may be accounted for with this methodology. So far there has been no direct way to measure representational change within human participants. However, Karmiloff-Smith [1992] mentions that computational modeling may provide important insights into specific brain processes that cannot be measured through other experimental means.

The theory of neuroconstructivism may provide some assistance with quantifying representations. A central assertion of this theory is that cognition affects the development of the brain, and the current state of the brain constrains cognitive development [Mareschal, et al. 2007; Polk & Farah 1998; Sirois, et al. 2008; Westermann, et al. 2007]. It is because of this interdependent relationship that any theory regarding cognitive function should be informed by underlying biological activity without reducing all

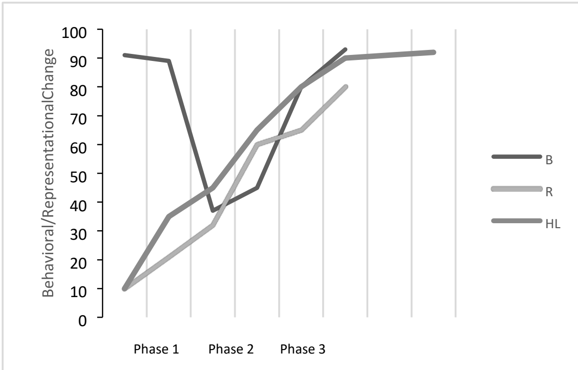

Figure 2: Conceptual comparison of a Neuroconstructivist model from Westermann and Ruh (2012) and the U-shaped curve, as originally presented by Karmiloff-Smith. Hidden layer development (HL) is comparable to representational development.

cognition to simply biological function. By investigating neural function, through the use of computational modeling, and the relationships between levels of description, such as behavior in relation to the state of a neural model, it might be possible to observe representational change in a more direct manner.

Various computational models have been developed with the intent to simulate neural function in a way that mimics observed human behavior, and these models might then provide insights into human neural functioning. One such model is used in Westermann [1998] to simulate brain activity in a way that closely approximates human learning of English past tense verbs. The model consists of a large number of nodes acting as an idealized version of a neural network. Nodes are assembled into an input layer, responding to a specific stimulus, a hidden layer that modifies the outputs of the model, and an output layer. In the neuroconstructivist model, the number of nodes present in the model is modified, in addition to the weights of the connections between nodes. In this way, the model modifies its structure by constantly evaluating each node in the hidden layer, and each of these nodes are replaced if they are found to be contributing to a large amount of errors according to certain algorithms operating within the system. Because each node in the hidden layer only responds to certain inputs, each node has a certain "receptive" field for which the node will activate, and modification of the hidden layer alters these fields. In this way, the model learns how to reproduce a large amount of past tense verbs when provided with the present tense version, and this is accomplished by changing the structure of the network in order to more effectively represent information.

After exposing the model to many different verbs, both regular and irregular, different trends about the model and its performance can be observed. In Westermann [1998], it was discovered that the model, which starts with only two nodes in the hidden layer, develops receptive fields that may overlap, implying that some nodes may respond to the same input as other nodes; however, there may be many inputs that are exclusive to either node. The hidden layer is described as a type of "memory" by the author, which responds to the identity of certain verbs rather than just the structural qualities of the word.

Additionally, Westermann and Ruh [2012] continued experimentation with this model and past tense acquisition, discovering that as the model was exposed to more verbs, the nodes in the hidden layer became more specialized to either regular or irregular verb forms. In addition, this model demonstrated U-shaped performance, i.e. failing to reproduce verbs that had been produced accurately at an earlier time. Model performance was then compared to results of previous studies involving human participants completing similar tasks, and a significant correlation was reported [Westermann & Ruh 2012].

The specific properties of this model seem to have implications for understanding cognitive development within humans. First, the late-occurring errors observed in both the model and human participants seem to indicate that there are similar processes taking place in both contexts. The hidden layer is also of interest in that the structural change to the model may indicate a change in how the model represents the information presented to it. In this way, the number of hidden units, the degree of specialization among these units, and the receptive field that each node possesses all seem to be direct indications of representational change and development at the level of a neural network. If this is the case, then there appears to be evidence for an increase in the complexity of representations across developmental time in the form of increased nodes in the hidden layer of the model (see Figure 2 for a visual depiction of this concept, as originally presented by Karmiloff-Smith which allows room for interpretation). The hidden layer also starts with very broad receptive fields (two nodes accounting for all inputs) that become increasingly complex in order to compensate for ineffective performance of the current network. This seems to mirror representational redescription on the cognitive level in response to exceptions to current theories.

By developing a computational model with similar behavioral qualities as human participants, inferences into how the brain and mind might be developing may be formed. However, this specific neuroconstructivist model only operates within the domain of English past-tense verb acquisition. In order to provide further insights into the other domains previously mentioned in this paper, it would be necessary to construct models that can operate within these domains. It may also be the case that such modeling can be applied to other fields of study, particularly those that utilize more advanced neural network models, such as artificial intelligence.

4 APPLICATIONS TO ARTIFICIAL INTELLIGENCE

Given the current popularity of brain-inspired computational representations (e.g., deep neural networks and deep reinforcement learners), knowledge about how representations develop and change in human brains is more relevant than ever. Taken together, the neuroconstructivist model and representational redescription theory offer a unique approach for the design and study of artificial intelligence. As demonstrated in Westermann and Ruh's [2012] model, representational change may provide a way for a machine to develop effective representations when presented with novel stimuli. This may provide a significant advantage over fixed programming, in that adaptation to open environments would become possible. By changing how information is stored, artificially intelligent agents may become better able to adapt to their environment, leading to more efficient automation with broader, more human-like capabilities.

Conversely, artificial intelligence research might also provide insights into the representational redescription theory that the current data with human participants has not yet been able to address. In the U-shaped curve, for instance, one of the main markers of the increase in the level of representation in phase 2 is a decrease in developmental performance. While this decrease in performance is adequate behavioral evidence for the curve, it does not sufficiently support the claim of representational change. Utilizing AI models in testing the existence of the U-shaped curve would allow for a clear, measurable, and observable separation between performance and the representations that facilitate behavior.

Other critical aspects must be considered when using artificially intelligent agents to observe representational change within different domains. As mentioned above, the neuroconstructivist model is only designed for learning past-tense forms of English verbs. Using neural models that can alter connection weights in addition to the structure of the network itself in the context of agents that can manipulate physical objects may assist in explaining the differences observed between studies investigating children's performance on the block balancing task. In fact, many other domains and tasks may be investigated in this manner.

5 CONCLUSIONS

The confusion on the meaning of the three phases and the ambiguity of representational redescription's central claims have led to a variety of incompatible interpretations, particularly with respect to the correct meaning and use of the U-shaped curve. There is a lack of generalized consensus among previous studies about what constitutes representational change in the context of this theory and finding means to test for this change that do not rely on the behavioral component of the U-shaped curve. The majority of prior studies did not operationalize representational change effectively as they based representational change solely off of observable behaviors. In this paper, we hoped to decrease the confusion about representation in Representational Redescription Theory. By examining neuroconstructivism in relation to representational redescription, representational change can be quantified without dependence on behavior as this theory states that cognition affects brain development; as knowledge becomes redescribed, the number of nodes increase, the nodes in the hidden layer become more specialized, and each node gains a receptive field, indicative of representational change. Thus, the theory of neuroconstructivism fills in the gaps in the central claims of representational redescription, as well as decreasing the confusion over what a representation is.

In this paper, we have surveyed the literature on representational redescription theory and the closely related neuroconstructivist theory, finding that they both (when properly interpreted) predict the existence of a U-shaped curve in the development of mental representations. We claim that this striking fact has implications for AI research, especially as the field continues to be dominated by increasingly massive neural-like networks. Namely, computational representations that claim to be realistic models of the mind must, at a minimum, be able to show how representations and performance match the predictions set out by representational redescription theory. In future work, we hope to expand on these implications for AI research.

REFERENCES

[1] Cheung, C., & Wong, W. (2011). Understanding conceptual development along the implicit–explicit dimension: Looking through the lens of the representational redescription model. Child Development, 82(6), 2037-2052. doi:10.1111/j.1467-8624.2011.01657.x

[2] Critten, S., Sheriston, L., & Mann, F. (2016). Young children's spelling representations and spelling strategies. Learning & Instruction, 4634-44.

[3] Critten, S., Pine, K. J., & Messer, D. J. (2013). Revealing children's implicit spelling representations. British Journal Of Developmental Psychology, 31(2), 198-211. doi:10.1111/bjdp.12000

[4] Critten, S., Pine, K., & Steffler, D. (2007). Spelling development in young children: A case of representational redescription?. Journal Of Educational Psychology, 99(1), 207-220. doi:10.1037/0022-0663.99.1.207

[5] Hollis, S., & Low, J. (2005). Karmiloff-Smith's RRM distinction between adjunctions and redescriptions: It's about time (and children's drawings). British Journal Of Developmental Psychology, 23(4), 623-644. doi:10.1348/026151005X35390

[6] Krist, H., Horz, H., & Schönfeld, T. (2005). Children's block balancing revisited: No evidence for representational redescription. Swiss Journal Of Psychology / Schweizerische Zeitschrift Für Psychologie / Revue Suisse De Psychologie, 64(3), 183-193. doi:10.1024/1421-0185.64.3.183

[7] Lorandi, A., & Karmiloff-Smith, A. (2012). From sensitivity to awareness: morphological knowledge and the Representational Redescription model. Letras De Hoje, 47(1), 6-16.

[8] Mareschal, D., Johnson, M. H., Sirois, S., Spratling, M. W., Thomas, M.S.C., & Westermann, G. (2007). Neuroconstructivism: How the brain constructs cognition. Oxford, New York: Oxford University Press.

[9] Messer, D. J., Pine, K. J., & Butler, C. (2008). Children's Behaviour and Cognitions across Different Balance Tasks. Learning And Instruction, 18(1), 42-53

[10] Patkowski, M. (2013). The critical period and parameter setting in five cases of delayed L1 acquisition. EUROSLA Yearbook, 13, 1–21. https://doi.org/10.1075/eurosla.13.03pat

[11] Picard, D., & Vinter, A. (2006). Decomposing and connecting object representations in 5- to 9-year-old children's drawing behaviour. British Journal Of Developmental Psychology, 24(3), 529-545. doi:10.1348/026151005X49836

[12] Pinker, S. (1995). *The language instinct: The new science of language and mind* (Vol. 7529). Penguin UK.

[13] Polk, T. A., & Farah, M. J. (1998). The neural development and organization of letter recognition: Evidence from functional neuroimaging, computational modeling, and behavioral studies. Proceedings of the National Academy of Sciences, 95(3), 847–852.

[14] Simpson, A., & Stehlíková, N. (2006). Apprehending Mathematical Structure: A Case Study of Coming to Understand a Commutative Ring. Educational Studies In Mathematics, 61(3), 347-371.

[15] Sirois, S., Spratling, M., Thomas, M. S. C., Westermann, G., Mareschal, D., & Johnson, M. H. (2008). Précis of Neuroconstructivism: How the Brain Constructs Cognition. Behavioral and Brain Sciences, 31(3). https://doi.org/10.1017/S0140525X0800407X

[16] Voutsina, C. (2012). Procedural and conceptual changes in young children's problem solving. Educational Studies In Mathematics, 79(2), 193-214. doi:10.1007/s10649-011-9334-1

[17] Westermann, G. (1998). Emergent modularity and U-shaped learning in a Constructivist neural network learning the English past tense. Proceedings of the 20th Annual Conference of the Cognitive Science Society, 1998.

[18] Westermann, G., Mareschal, D., Johnson, M. H., Sirois, S., Spratling, M. W., & Thomas, M. S. C. (2007). Neuroconstructivism. Developmental Science, 10(1), 75–83. https://doi.org/10.1111/j.1467-7687.2007.00567.x

[19] Westermann, G., & Ruh, N. (2012). A Neuroconstructivist Model of Past Tense Development and Processing. Psychological Review, 119(3), 649–667. https://doi.org/10.1037/a0028258

[20] Xie, K., Fox, G. E., Liu, J., Lyu, C., Lee, J. C., Kuang, H., … Tsien, J. Z. (2016). Brain Computation Is Organized via Power-of-Two-Based Permutation Logic. Frontiers in Systems Neuroscience, 10. https://doi.org/10.3389/fnsys.2016.00095